



Moral circle expansion: A promising strategy to impact the far future

Jacy Reese Anthis^{a,b,*}, Eze Paez^{c,d}

^a The University of Chicago, USA

^b Sentience Institute, USA

^c Pompeu Fabra University, Spain

^d Centre de Recherche en Éthique, Université de Montréal, Canada

ARTICLE INFO

Keywords:

Existential risk
Effective altruism
Artificial intelligence
Animal ethics
Longtermism

ABSTRACT

Many sentient beings suffer serious harms due to a lack of moral consideration. Importantly, such harms could also occur to a potentially astronomical number of morally considerable future beings. This paper argues that, to prevent such existential risks, we should prioritise the strategy of expanding humanity's moral circle to include, ideally, all sentient beings. We present empirical evidence that, at micro- and macro-levels of society, increased concern for members of some outlying groups facilitates concern for others. We argue that the perspective of moral circle expansion can reveal and clarify important issues in futures studies, particularly regarding animal ethics and artificial intelligence. While the case for moral circle expansion does not hinge on specific moral criteria, we focus on sentience as the most recommendable policy when deciding, as we do, under moral uncertainty. We also address various nuances of adjusting the moral circle, such as the risk of over-expansion.

1. Introduction

There are currently around 8 billion humans (10^9) on Earth. There are over 100 billion domestic animals (10^{11})—primarily chickens and fishes used in the food industry. There may be trillions of wild birds and mammals (10^{12}) and over a quintillion tiny wild animals such as insects (10^{18}) (Tomasik, 2009). These moral stakes are difficult to conceptualise, yet they pale in comparison to the potential long-term scope of human civilisation in the distant future, such as with interstellar expansion.¹

Some moral philosophers and futures scholars are beginning to consider our impact on the far future as a serious and perhaps overwhelmingly important consideration that we should account for in present decisions (see, for example, Beard, Rowe, & James, 2020; Bostrom, 2014; Liu, Lauta, & Maas, 2018; Kareiva & Carranza, 2018; Moynihan, 2020; Parfit, 1992). Such work is part of a broader project of prescriptive futures studies, asking questions such as how to encourage thoughtful contemplation of the future, how to design future-oriented public policy, and how to account for the indirect effects of our actions on the future (Ahvenharju,

* Corresponding author.

E-mail addresses: jacy@uchicago.edu (J.R. Anthis), joseezequiel.paez@upf.edu (E. Paez).

¹ For example, suppose humanity expanded to the Virgo Supercluster. Then the human population could be one hundred undecillion (10^{38}). See Bostrom (2003). Notice, however, that the same interstellar resources could fuel many more minds of less complexity, and therefore less energy cost, than the human mind. Thus, assuming minds less complex than the human mind can be sentient, the total number of possible future sentient individuals is proportionally higher.

Minkinen, & Lalot, 2018; Dorsser, Walker, Taneja, & Marchau, 2018; Ng, 2020). Restricting our concern to present or near-future generations of human beings may be a significant normative blindspot, neglecting the majority of individuals who will, at any time, exist. While work on 'existential risk' has focused on extinction risk in particular, recent work has highlighted the importance of 'suffering risks', which are 'astronomical in scope and hellish in severity' (Kovic, 2021). Human history is riddled with atrocities committed against contemporary members of our own species, and if species-based discrimination is wrong, then history is also riddled with serious harms committed against nonhuman animals. Thus, the risk that we commit atrocities on an interstellar scale is a serious danger—perhaps one of the greatest dangers that exist on the horizon of our species.

In this article we present a novel approach to reducing such existential risks and bettering the future. We build on an interdisciplinary literature including animal ethics, environmental ethics, ethical theory, history, and social psychology. We defend two claims. The first is that to minimise the danger of existential risks, particularly suffering risks, we should prioritise a strategy that we call:

Moral circle expansion: a community's moral circle has expanded with respect to some previous time t if, and only if, a number of entities which used to be given less than full moral consideration at t are now given more moral consideration.

Clearly laying out this definition allows us to describe the following familiar phenomenon as a case of moral circle expansion: a community of moral agents concludes that a feature on which they previously based the moral inconsiderability of some individuals is not morally relevant. Thus, they extend moral consideration to all members of the set of entities that possess that feature. This definition is completely agnostic about what attributes—sentience, rationality, being alive, etc.—a community uses to construct its moral circle. In addition, it is not a moralised definition. Thus, according to our definition, in order to identify some shift in a community's sphere of concern as an instance of moral circle expansion, it is not necessary to judge whether they are deploying a more justified criteria of moral considerability. This has several advantages. First, we do not need to endorse the same criteria in order to agree that a community's moral circle has expanded. Second, our definition allows us to identify instances of moral circle expansion in a society and retain the ability to criticise them (i.e., one can believe a specific instance of moral circle expansion was a mistake).

Note, finally, that this is a definition of, so to speak, gross expansions of the moral circle. It could be modified to accommodate calculations of *net* expansion. This would consist of something like the number of new entities included with respect to t minus the number of entities excluded with respect to t , given an appropriate specification of 'minus' that clarifies the subtractivity of 'sets'.²

We will argue that explicitly discussing and analysing moral circle expansion is a revealing yet underexplored perspective for understanding the nature of moral progress. While much scholarly work can fit under the umbrella of moral inclusion and exclusion (e.g., prejudice, discrimination), anything more than a passing reference to the moral circle is rare, and we argue there is much to be gained by direct exploration of moral circle expansion itself.

Our second claim is that humanity's moral circle ought to be expanded to include, ideally, all sentient beings. That is, all beings with a capacity for positive and negative experiences.³ In this paper, we will not assume that sentience is necessary for moral considerability. However, looking back on historical atrocities, it seems that many could have been prevented or mitigated if communities had extended their moral concern to other groups of sentient beings, such as humans of various races, ethnicities, sexualities, genders, or nationalities—or animals of various species who share this planet with humankind (Judge & Wilson, 2015; Wright, 2018). Thus, even if one is unsure what exactly the future of the moral circle should look like (e.g., the moral patienthood of artificial intelligence), pushing on the current frontiers of the moral circle (e.g., farmed animals) or otherwise engendering expansion towards other kinds of sentient beings is a compelling moral priority not only for utilitarians and others concerned with doing the most good but, more generally, for those concerned with preventing serious wrongs.

Artificial intelligence researcher Paul Christiano (2013) argues that, for those trying to do the most good, 'changing long-term social values' should not be a priority for several reasons, particularly the changing nature of moral values upon reflection and the potentially zero-sum nature of competing moral values. In addition, there are many other pressing issues demanding our limited attention, such as the reduction of extinction risk (Bostrom, 2003), the mitigation of the effects of anthropogenic climate change, and the alleviation of global poverty.⁴ Thus, the claim that moral circle expansion should be prioritised, among all of those compelling projects, is in need of some defence.

² As we mentioned, a society's expansion of its moral circle usually occurs in terms of sets, rather than individuals. That is, increased moral consideration is usually given to members of the set of individuals who possess a particular attribute. This is consistent with prominent accounts of discrimination in terms of the exclusion of societal groups (see, for example, Lippert-Rasmussen, 2013). We decided to define gross moral circle expansion in terms of individuals, rather than in terms of set, because of the difficulties the latter generates. First, for a given situation, there may be various equally acceptable criteria for carving up sets and there seems to be no value-neutral way to decide which to choose. Employing different criteria, the same societal shift may be described either as an expansion or a contraction. Second, even if there were reasons to prefer some set constructions over others, it would be necessary to specify rules for the subtraction of sets when calculating net expansion. For example, suppose a society stops giving moral consideration to members of set A (1000 individuals), but expands its circle to include members of sets B and C (100 individuals each). A simple additive rule would count that as an instance of moral circle expansion, which may be counterintuitive, given the net exclusion of individuals.

³ We are aware that, on some views, non-sentient individuals (e.g., non-sentient living organisms, biocenoses, ecosystems, species, biodiversity, or landscapes) also merit moral consideration and would be legitimate objects of moral circle expansion. We are also aware that, on some other views, individuals with sophisticated cognitive capacities merit greater moral consideration than merely sentient ones. Though we are not persuaded by these views, our defence of moral circle expansion is compatible with them. We elaborate further on this point in Section 3.2.

⁴ In theory, these moral issues could be approached via moral circle expansion for future beings, the natural environment, or the global poor, but in practice other approaches are more common.

Indeed, most laypeople today do not appear to treat moral circle expansion as a moral priority. For example, while Americans care much about their dogs and cats at home, around 99 % of U.S. farmed animals today live on factory farms, where they are subject to intense confinement and suffering (J. R. Anthis, 2019). Around 87 % of U.S. adults believe ‘farmed animals have roughly the same ability to feel pain and discomfort as humans’ (Reese, 2017). Yet, U.S. rates of self-identified vegetarianism have been stable around 5–6 % since 1999 (Reinhart, 2018). While there may be redeeming bits of moral reasoning at play, such as certain views of personal responsibility, this is a *prima facie* case that the general human population is not prioritising moral circle expansion nearly as much as they could.

We proceed as follows. First, we elaborate on the concept of moral circle expansion and outline the ethical case for it, as well as its empirical underpinnings and relevant discussion in ethics and philosophy. Second, we address the topic of moral uncertainty, arguing for expansion to all sentient beings and addressing the risk of over-expansion. Third, we explore what the future of moral circle expansion may look like, both in the abstract and in practice. Finally, we return to the topic of the far future, arguing that moral circle expansion may be a useful strategy to reduce the likelihood of plausible astronomically adverse outcomes. We end with some practical remarks on the role of moral circle expansion in philosophy, social science, and futures studies.

2. Moral circle expansion: ethics, history, and psychology

2.1. The ethical case for expanding the moral circle

The concept of varying moral concern for different entities is fairly intuitive and was widely discussed by philosophers throughout history, but the first modern use of the ‘circle’ analogy and the first discussion of an expanding circle is attributed to historian William Edward Hartpole Lecky (1869).⁵ He discussed why people of different time periods would use different moral perspectives to evaluate the same event, such as human gladiators killing each other for sport. He argued that while humans have differed greatly in their specific moral standards, there has been a ‘natural history’ that has moved over time in the direction of including more individuals in the ‘circle’ of ‘benevolent affections’. The ‘circle’ analogy was popularised by utilitarian philosopher Peter Singer (1981), who argued that human altruism began with helping family members but is now buoyed by rational thought to include a wide range of beings in an expanding moral circle.

The ethical case for expanding humanity’s moral circle to all sentient beings is fairly robust. It relies on merely two plausible claims. The first is the principle that all sentient beings matter, at least to some extent. A standard defence of this claim is the so-called ‘argument from relevance’ (Horta, 2018): moral reasoning consists, in part, of identifying the reasons we have as agents for or against the alternative courses of action available to us. One source of such reasons is the impact of our choices, either positive or negative, on all entities who might be affected by them. A sufficient possibility condition for an entity to be affected is to be sentient, that is, to possess a capacity for phenomenal states of positive or negative valence. Thus, when engaging in practical deliberation, agents should consider the effects of their choices on all sentient beings who can be benefited or harmed by them. The second claim is that humans currently consider the interests of less than all sentient beings. For example, most humans routinely disconsider the interests of nonhuman animals, despite leading scientists considering many of them to be sentient (Low, 2012).

Importantly, this argument for expanding our moral circle makes no other normative or axiological assumptions. It is thus compatible with different normative views (e.g., consequentialism, deontology) and different values such as happiness, freedom, and justice. Furthermore, it is also compatible with the claim that other, non-sentient entities may matter as well. It is even possible to combine it with a hierarchical approach regarding moral considerability. On this view, though all sentient beings matter, some do so more than others (Kagan, 2019).

The delineation of who receives moral concern and who does not—whether based on sentience or other criteria—is rarely a tidy circle centred on the self. It is more like a multidimensional gradient (e.g., physical proximity, social proximity, similarity to the self) of concern that can range from malicious intent for beings whom one would prefer to be harmed to extreme altruism for beings whom one cares about more than oneself. There are also different ways to measure the moral circle. We could pay attention to the *self-reported moral circle*, or the beliefs about who matters, and attendant desires, as expressed by a moral agent. We could also pay attention to the *choice-based moral circle*, or the beliefs about who matters, and attendant desires, as implied by a moral agent’s choices. We can also think of these as an *attitudinal moral circle* and a *behavioural moral circle*. There is of course a spectrum between attitudes and behaviour, such as how people vote for new laws or politicians, which is expressing an attitude, but one with higher stakes than just answering a survey question.

Moral circle expansion has been located within the emerging scholarly literature on *effective altruism*, using evidence and reason to help others as much as possible (see, for example, Mohorčič & Anthis, 2018; Reese, 2018; Singer, 2015; Smith & Reese, 2016), and it aligns with the recent focus on *longtermism*, the prioritisation of long-term impact (Greaves & MacAskill, 2019; Ord, 2020). Similarly, moral circle expansion could be a means of addressing ‘ongoing moral catastrophes’, that is, very serious moral wrongs which, unbeknownst to us, we presently do (Williams, 2015).

⁵ There is an earlier use of ‘circle’ in the work of Hierocles, a 2nd century Stoic philosopher. Records of Hierocles’ work are fragmented, but he discusses the social circles of an individual person and their duties towards those around them—stronger towards those closest—in *On Appropriate Acts*, though he does not discuss the expansion of the circle over time (Ramelli, 2009).

2.2. Historical evidence for moral circle expansion

The usefulness of moral circle expansion depends on its reliability. Namely, does the moral circle exist in the sense that altruists can take action that will, in expectation, lead to more groups being morally considered? We can provide a *prima facie* validation of this meliorism in the historical record. For example, can working on the moral issues of the day in which one lives, such as an altruist in the 1800s fighting for the abolition of the transatlantic slave trade, beget further moral circle expansion, such as animal rights in the 2000s? Does a rising moral tide eventually lift all boats?

We can separate this expansion process into two distinct, intuitively plausible subprocesses. First, an expanded moral circle makes traction on the direct issue (e.g., moral concern for slaves leading to the mitigation or abolition of slavery). Second, that bit of moral progress facilitates future moral progress (e.g., the morally driven abolition of slavery engenders greater concern for animals centuries later).

Regarding the first, it seems fairly straightforward *prima facie* that moral consideration of the victims of atrocities could have prevented or mitigated them, since atrocities necessarily cut against the interests of their victims. If pre-nineteenth century Europeans had more fully considered the interests of Africans, then the transatlantic slave trade could have been abolished much sooner. If people today cared more about animals, then presumably they would not exploit or abuse them, as long as those people had the material means to do so.

Regarding the second step in moral circle expansion, looking back at history also shows us a number of people who seem to have peered ahead on the moral arc, prefiguratively connecting the dots between moral issues with remarkable foresight. Consider, for example, Jeremy Bentham and John Stuart Mill in the eighteenth and nineteenth centuries. They advocated for abolition of slavery, equality of men and women, decriminalisation of homosexuality, personal liberty, and the moral consideration of animals (Bentham & Crompton, 1978; Bentham, 1789; Mill, 1859).

How did they do this? Bentham and Mill are best known for the development and propagation of utilitarianism. Of course, one need not be a utilitarian to have the ability to peer ahead on the moral arc, as we say. Basing our choices in how our actions affect other individuals' lives, by their own lights, and being as ecumenical as possible regarding who we include in our moral deliberation seems a sound way to minimise the risk of committing some serious moral wrong. Consider, for example, how some of the most prominent antislavery advocates such as William Wilberforce and George Thorndike Angell were also founders of some of the first animal rights groups such as the Royal Society for the Prevention of Cruelty to Animals and the Massachusetts Society for the Prevention of Cruelty to Animals (Belmonte, 2007; MSPCA, 2015).

But perhaps moral inclusion is not sparking further moral inclusion, and instead further moral inclusion merely correlates with but is exclusively or mostly caused by an external variable such as economic progress or technological efficiency. For example, if you use material wealth, such as Gross Domestic Product, as a proxy for human wellbeing, then history looks like a hockey-stick graph, flat for thousands of years with a near-vertical leap in the mid-1800s coinciding with the industrial revolution. Historical causation is notoriously challenging to disentangle, but there has been a moderate amount of debate on moral versus economic explanations for the decline and abolition of the transatlantic slave trade—primarily the British trade, which was outlawed by the UK Parliament in 1807. In this example, the evidence overall suggests that abolition resulted primarily from a social movement that invoked moral concerns rather than economic ones, and that in fact, the transatlantic slave trade was highly profitable in the years leading up to its abolition (Drescher, 1977).

Over the past decade, there has been an emergence of progress literature—articles and books that track the expansion of humanity's moral circle over the past few centuries and more broadly detail how life for humans has gotten substantially better despite what the public hears in the news (see, for example, Easterbrook, 2018; Pinker, 2018).

There is evidence that moral circle expansion in the human context is now spilling over into increased concern for nonhuman animals, as outlined in Reese (2018). The past few decades have seen a remarkable shift in scientific opinion towards animal sentience. Animal behaviourist Donald Griffin (1981) received widespread dismissal for even addressing the topic of animal awareness, despite his prominence following his monumental discovery of echolocation in bats. However, in 2012, the Cambridge Declaration on Consciousness, written by a group of prominent neuroscientists, asserted that 'the weight of evidence indicates that humans are not unique in possessing the neurological substrates that generate consciousness' (Low, 2012).

We should qualify here that while it makes sense to consider an expansion of the moral circle to nonhuman animals as a current frontier, we do not expect the expansion to occur all-at-once for the animal kingdom, or even any large swath of it such as primates. Expansion can be very uneven. Consider, within the scope of animals used for food, how there is already widespread moral repulsion in the West at eating dog meat, yet in the case of fish, very dissimilar beings, there is hardly any discussion of their welfare or mainstream efforts to consider their interests in any way. Expansion may progress differently for different animals based on their physical features, their use in human society, and where they fit into the 'mixed communities' in which humans and animals coexist (Midgley, 1983). Similar factors may predict unevenness in the expansion of our concern for non-sentient organisms such as plants, collective entities such as ecosystems, and artificial entities such as robots—in each category, one can imagine entities for which expansion will be easier or harder.

None of this implies, however, that future moral circle expansion is inevitable. Indeed, there are several ways in which the moral circle seems to have narrowed from one period of time to another. Some societies assign now less importance to the preferences of deceased family members and ancestors, or care less about supernatural entities and the natural environment (Reese, 2018, 16). As Rodman (1979) observes, some ancient Roman philosophers included 'all animals' in discussions of natural rights, while modern philosophers favoured a more restrictive view of the 'sphere of justice', which excluded nonhuman animals.

If we consider contemporary attitudes towards animal well-being, the self-reported moral circle appears to be expanding in some

societies. However, it seems that the choice-based moral circle has clearly narrowed if we include the welfare of the over 100 billion animals living on factory farms (K. Anthis & J. R. Anthis, 2019).⁶

2.3. The psychological underpinnings of the moral circle

We can also look to the psychology literature to validate moral circle expansion as a tangible phenomenon. For instance, in the case of the contemporary inclusion of animals, there is already a nascent literature on the psychology of *speciesism*, or species-based discrimination, which is analogised to racism, sexism, classism, and other human prejudices. Namely, speciesism has been introduced as ‘a measurable, stable construct with high interpersonal differences, that goes along with a cluster of other forms of prejudice, and is able to predict real-world decision-making and behaviour’ (Caviola, Everett, & Faber, 2018).

Indeed, there is a growing area of research exploring the correlations between particular forms of prejudice, such as racism, sexism, or speciesism, and other more general psychological dispositions, such as social dominance orientation (SDO) (Sidanius & Pratto, 1999, 61). There is some evidence that SDO is a disposition also connected to racial and species prejudice (Caviola et al., 2018; Dhont, Hodson, Costello, & MacInnis, 2014; Dhont, Hodson, & Leite, 2016; Graça, Calheiros, Oliveira, & Milfont, 2018). To a lesser extent, it seems that political conservatism is relevant to speciesist attitudes, and right-wing authoritarianism seems to be connected to the perception of vegetarianism as a threat (Caviola et al., 2018; Dhont & Hodson, 2014).

Furthermore, Hodson, Macinnis, and Costello (2013) suggest that the human-animal hierarchical division is related to the belief that some human groups merit less moral consideration. Their Interspecies Model of Prejudice proposes that hierarchical beliefs towards animals correlate with greater hostility towards human outgroups. These beliefs operate by systematically reinforcing the negative connotation of animal dehumanization and highlighting those attributes which are perceived as making these humans more similar to nonhuman animals (Costello & Hodson, 2010; Hodson et al., 2013). Moreover, a belief in human-animal hierarchical division predicts greater racial prejudice both in children and in adults (Costello & Hodson, 2014).

There is also a small yet compelling literature on ‘moral expansiveness’, in particular the development of a Moral Expansiveness Scale (MES), which uniquely predicts willingness to help others at the cost of personal and in-group self-interest. Current results suggest this approximate ordering: family/friends, in-group, revered, stigmatised, out-group, animals (high sentience), environment, animals (low sentience), plants, and villains (Crimston, Bain, Hornsey, & Bastian, 2016).⁷

Overall these psychological findings suggest the existence of a tangible individual moral circle, such that increased concern for a specific population may be interlinked with concern for other specific populations. Given the breadth of issues that a rigorous understanding of moral circle expansion would encompass, and, given the initial evidence for the phenomenon from history and psychology, it seems further research is warranted.

3. The problem of moral uncertainty

Moral circle expansion is an instance of a broader phenomenon, which we may call *moral circle adjustment*. It occurs when a community modifies its attitudes and choices in order to keep them coherent with the account of moral considerability that it believes to be, all things considered, more justified. However, just as a community’s moral circle can be adjusted to assign more moral considerability to certain individuals, it can be changed to assign them less or to exclude some individuals altogether from the moral circle (i.e., moral circle narrowing).

If moral claims are simply an expression of our attitudes, or our preferences, we need not be concerned that, in adjusting our moral circle, we may have failed to grasp important subject-independent moral truths. But if we believe that there are such truths, it is possible to adjust the moral circle in the wrong way. This will happen if we base our adjustment on an account of moral considerability that we were mistaken to endorse. Furthermore, we may not assign absolute certainty to our preferred moral theories. We may believe there is some probability that some other account of moral considerability is the true or most justified one. Indeed, because these are complex problems, it may not be justified for us to give *full* credence to our preferred views, even if they are the ones to which we give the *highest* credence. Thus, even when we adjust our moral circle according to what we believe to be the best moral theories, we may be concerned that, unbeknownst to us, we are committing serious wrongs.

For instance, when adjusting our moral circle, there is the risk that we over-expand. Suppose that a community has some pool of limited resources. At first, their moral circle is based on the belief that only individuals with some property, *P*, deserve moral consideration and thus have a claim to those resources. However, the community no longer believes that *P* is necessary for moral considerability, so they extend their concern to a much larger group of individuals, some of whom do not have *P*, and distribute their

⁶ If moral circle expansion is not inevitable, then those trying to do the most good can change the *direction* of moral change, which may do much more good than simply increasing its *speed*. For example, technological progress tends to be monotonic, and thus inevitable, at least without major social catastrophe such as nuclear war, while social progress trends up and down at different times. However, speeding up expansion may change direction in subtle ways, such as affecting humanity’s values at the occurrence of *value lock-in* (e.g., because from that point on, an artificial intelligence is the primary decision-maker in human society) or encouraging change to happen during a *golden era* (e.g., because during this time, technological and social factors make moral progress easier than it would be at a later time). See Anthis (2018) for more discussion.

⁷ The low ranking of ‘villains’ has potentially important implications for altruistic behaviour. Depending on one’s view of moral desert, whether wrongdoing makes one’s welfare actually matter less, this could make advocacy on behalf of convicted wrongdoers—such as criminal justice reform—a particularly important cause area within moral circle expansion.

resources accordingly. This means that the average included individual has, all else equal, a much smaller share of resources. Because of this, members of the previously privileged group are much worse off than they would have been had their community retained their old moral beliefs. In this example, suppose that the original belief, that only individuals with *P* deserve moral consideration, was actually the most justified one. By depriving these individuals of resources to which they had a claim, a serious moral wrong has been committed. Since we can be mistaken in this way, it could be said, moral circle expansion should not be a priority.

We do not think this objection works. First, moral circle expansion should be pursued only when recommended by the decision procedure we ought to employ when choosing under axiological and normative uncertainty. Let us consider the *choice-worthiness* of a policy as its desirability according to some theory. Some suggest that the choice-worthiness of an option should be weighted according to the credence we give to the theory that recommends it (MacAskill & Ord, 2018). This would give us a measure of its *expected* choice-worthiness. In cases in which our uncertainty ranges across theories that (a) quantify the choice-worthiness of options and (b) allow for inter-theoretical comparisons of choice-worthiness, our aim should be, precisely, to maximise expected choice-worthiness. Thus, for these kind of theories, maximising choice-worthiness would be the right decision procedure under normative (moral or prudential) uncertainty. This mirrors the role expected utility theory has as the decision procedure under epistemic uncertainty. This does not immunise us from making mistakes. Since, however, we would have done all we rationally can in order to choose the morally right option, we would not be blameworthy.

Second, expanding our moral circle to include all sentient beings is recommended by this approach. Note that we are not advancing the claim that sentience is sufficient, or necessary, for full moral considerability. We believe that sentience suffices to render an individual morally considerable in a non-trivial way. We also believe this entails that we have strong reasons not to inflict them serious harms. These are claims many others may be ready to accept.

Some people believe that, in addition, possessing sophisticated cognitive capacities makes an individual more morally considerable. Some environmentalists believe, as well, perhaps more controversially, that we should extend moral consideration to non-sentient natural entities. Even though we do not necessarily endorse these views, they are compatible with the claims we defend here. On the plausible assumption that sentient beings matter, our collective decisions will have a greater chance to be morally justified if our moral circle includes them all, even if we believe that other entities are morally considerable or that some sentient beings (e.g., those with larger cognitive capacities) merit special consideration.

Finally, we think there are *a posteriori* reasons to expect mistakes of under-inclusion in the moral circle to be worse than mistakes of over-inclusion. There is evidence from humanity's historical record that, regarding our moral circle, false negatives (i.e., a lack of moral consideration for a morally considerable group) are far more prevalent than false positives. In addition, as we will see in the next section, there are many categories of sentient beings to which, plausibly, human communities do not give sufficient moral consideration. A strategy of moral circle expansion would, at least in part, compensate for that trend.

4. Beyond the frontiers of our present moral circle

Just as Bentham and Mill peered ahead to some of the biggest social issues of the twenty-first century, an important starting point for encouraging moral circle expansion now and in the future may be to estimate some of the biggest moral issues of the coming centuries. To do this, we should examine the dimensions by which our moral circle has already expanded to a large extent (e.g., gender, race, physical and social proximity) and ask ourselves what similar dimensions there are by which our moral circle has yet to expand.

4.1. Wild animals and invertebrates

The fact that there are over 100 billion animals on factory farms is partly why we consider them one of the most important frontiers of today's moral circle (K. Anthis & J. R. Anthis, 2019). However, when it comes to sentient life on the planet, the vast majority of individuals are wild animals, on the order of a trillion wild birds and mammals and perhaps over a quintillion small invertebrates such as insects (Tomasik, 2009). Many contemporary humans care about wild animals to an extent, particularly charismatic megafauna such as pandas, elephants, and blue whales. However, the masses of wild animals, beyond those individuals, are seldom viewed as more than a natural resource, whether that is a beauty to be preserved or a tool for human benefit. The wellbeing of these animals, especially the suffering they endure from illness, starvation, and other natural causes, is rarely considered an ethical issue. Indeed, several authors believe that it is likely that suffering predominates in their lives because of these natural causes (Faria, 2016; Faria & Paez, 2015; Horta, 2010; Tomasik, 2015b).

Since wild animals constitute the majority of sentient individuals and it is possible that a great number of them have negative wellbeing, it seems particularly relevant to determine whether they could or should be the next frontier, after domestic animals, of the expanding moral circle. Just as humanity as a whole no longer thinks it is right to deny care for humans who suffer from natural causes—such as a natural disaster that strikes a foreign country—we might one day think it wrong not to address wild animal welfare by providing housing, food, and water to wild animals struck by those same natural disasters.

Despite the possible prevalence of suffering in nature, animal ethicists have not been overly concerned about our duties towards wild animals, with some exceptions. Some suggest, on the basis of a relational account, that it is morally permissible to alleviate or prevent wild animal suffering caused by natural events, but that there is no general moral requirement to do so (Palmer, 2010). Some utilitarian, telic egalitarian, and telic prioritarian proposals argue for an obligation to intervene in nature in order to help these animals (Pearce, 1995; Faria, 2014; Horta, 2010; McMahan, 2015). Such an obligation has also been advanced from rule-consequentialist (Paez, 2019), Kantian positions (Paez, 2020), liberal cosmopolitanism (Cochrane, 2018), the capabilities approach (Nussbaum, 2006, 325–407), and neorepublicanism (Paez, 2021). Because the natural status quo itself is often considered the main threat to animal

well-being, it has even been suggested that we should redesign nature and animal organisms in order to prevent or alleviate their suffering (Pearce, 1995; Paez, 2020).

Other authors suggest that wild animals must be regarded as forming their own sovereign communities in nature, and our relations with them should be governed by a general principle of non-intervention (Donaldson & Kymlicka, 2011; Korsgaard, 2018). In addition, the claim that we ought to intervene in nature rails against environmentalist intuitions. Some claim that natural wholes (Leopold, 1949) or living organisms (Taylor, 1981) possess inherent value. Preserving them is, on these views, one of our fundamental aims of action. But, as referenced above, evidence suggests that the ordinary workings of nature are likely to be seriously harmful for most sentient individuals. Therefore, in many cases, the aims of preserving nature and advancing the interests of sentient beings may be incompatible. One may only be achieved at the expense of the other. If a community's moral circle includes both wild animals and non-sentient nature, moral agents may be forced to choose between conservation and intervention in nature on behalf of animals.

Importantly, even if humanity ultimately decides on a non-interventionist stance, an expanded moral circle could make that decision more well-reasoned and considerate of wild animals' interests. This includes the interests of invertebrates with sufficient likelihood of sentience, such as insects or cephalopods. Given their extremely large numbers (Tomasik, 2009), they merit further empirical research on their capacity for suffering and, eventually, maybe their inclusion in our moral circle.

4.2. Future sentient beings

We have already implicitly discussed an astronomically large class of beings: those who might exist in the future, especially in the very long run of human existence. Human beings routinely display egotistic concern about events that might impact their future selves, as they plan for retirement or even just for the next weekend. We also regularly consider some beings who have not been born yet, such as when we plan for parenthood or worry about the effects of climate change on our children's children in the coming decades. But this is only the tip of the iceberg, given humanity's potential for continued existence over thousands of years, colonisation of other planets, or technological ability to create large numbers of sentient beings. We elaborate on this particular frontier of the moral circle in the following section.

5. The moral danger of the far future

So far, we have considered moral circle expansion from a social and an individual perspective and discussed a variety of factors involved in the expansion or narrowing of humanity's moral circle. We have also mentioned certain categories of individuals that an expanded moral circle ought to include. In this last section we will elaborate on the importance of effecting that inclusion from the perspective of the far future. How likely is it that a failure to sufficiently expand humanity's moral circle could lead to a far future that is much worse than it otherwise would be? We argue that the risk is high, given humanity's track record, risks of unwarranted optimism, and the possibility of artificial sentience.

5.1. Humanity's track record

Our treatment of animals is a particularly concerning example, especially in the food system. Humans crowd billions of chickens in dark, filthy sheds, where they suffer from extreme distress and the inhibition of their natural behaviours (HSUS, 2013). We breed them for ultra-fast growth to increase profits, which leads to animals collapsing from their sheer weight and dying because their internal organs cannot keep up with their muscle growth (NCC, 2019). We kill trillions of fish by decompression or suffocating them in air (HSA, 2018). Humans also decimate the natural environment. If this is what humans have done with our current power over just one planet, what might we do on a galactic scale?

It seems that significant backsliding, such as the moral circle no longer including a group for which significant strides have previously been made, such as women, people of colour, or children, is not very likely. Consider for example the forces described as driving moral progress in Pinker (2012): The Leviathan (a strong government with a monopoly on violence), Commerce, Feminization, Cosmopolitanism, and The Escalator of Reason. These forces seem likely to continue into the future (Reese, 2018). This continuation could require certain political or economic structures, and some catastrophic events, such as nuclear war or climate change, could fundamentally alter human society. However, it is not clear if this would lead to more or less of these forces. For example, a revived post-apocalyptic human society could face resource constraints coupled with technological power that encourages tribalism and authoritarian government, which may lead to narrower moral circles. On the other hand, a society that has endured a near-extinction event could have a greater sense of camaraderie due to that shared experience and the saliency of global-scale threats.

However, there may be a large possibility of stagnation, a moral circle that continues to expand but does not fully reach all sentient beings. Consider, for example, if humanity is one day able to create artificial sentience, such as sentient computer programs that exist only in the digital world (Buttazzo, 2001). These beings might be so unlike humans that we fail to care about them sufficiently or even recognise their sentience. This issue also arises with animals who are less physically similar to humans, such as how much easier it is to get people to care about dogs than pigs and fish.

5.2. Unwarranted optimism

There are reasons beyond our track record to be worried about the quality of the far future. One is the possibility that controversial axiological beliefs may lead us to be too optimistic in our value judgements. Consider the following thought experiment. Suppose you

are given a choice between creating some new world populated by beings similar to us or refusing to do so. You are then told that, if you choose the first option, there is a 99 % chance of a future in which each individual's life is as good as one can imagine, much better than that of the happiest human that has ever existed. Yet there is a 1 % chance that all future lives will instead be full of suffering, much worse than that of the worst-off human that has ever existed. What would be the morally justified choice?

Some people may be disposed to opt for creating this new world, given the contrast between 99 % and 1% and the intuitive comparability between a very good and very bad outcome. This intuitive judgment may be grounded in a belief such as that any unit of happiness has as much moral weight as any comparable unit of suffering, or that the maximal happiness of one individual cancels out the maximal suffering of another individual. Because in the original scenario, of any hundred possible futures, 99 contain lives full of happiness, the expected net value of choosing to create the new world may be seen as highly positive.

Yet many others may disagree. Consider what those two possible outcomes may actually look like. Certainly, in the utopia, any given person at any given time would enjoy the best possible experiences and find her life extremely fulfilling, such as with continuous experiences of the feeling after the birth of a parent's first child, an orgasm, or the best culinary experience at the world's finest restaurant. In the dystopia, however, life would perpetually feel much worse than being constantly burned alive, while watching the torture of everyone that individual loves, making it a veritable hellscape.

Once the outcomes are more vividly represented in this way, some may think that the 1% risk of this very bad outcome outweighs the 99 % chance of a utopia. Some may even claim that the risk would have to be as low as 0.1 % or 0.01 % to justify the choice for this new world. This could be grounded in a belief in a simple negative-leaning axiology, according to which any unit of suffering has more moral weight than any comparable unit of happiness.⁸ Consider a similar, but intrapersonal, scenario, in which you are asked how many days of bliss would compensate for you one full day of constant torture. If you believe that not even 100 days of bliss could compensate for that day of torture, then in the previous imaginary situation, you should probably refuse to create the new world. Furthermore, suppose someone endorses a negative axiology, believing that only suffering counts morally, having always negative value. Happiness counts for zero. For these people only a 0 % risk of dystopia could justify the choice for this new world.

If we reject a negative or negative-leaning axiology, we may have reasons, given our broader ethical view, to prefer the creation of the new world. But it also may be rationally important that intelligent people who have thought seriously about the issue disagree with us. Again, if there are subject-independent moral truths, it may not be justified for us to assign full credence to our views, even if they are the ones to which we give the highest credence. We may be mistaken. In this case, as mentioned, on some views, we should take into account how likely we think it is that the views we disagree with turn out to be the true or most justified ones (MacAskill & Ord, 2018). In this case the rational action is to apply a discount on our optimism grounded on the extent to which we find negative or negative-leaning views credible.

On this view, it may indeed turn out that a relatively small chance of dystopia would make the option of not creating that new world preferable. That would imply that our moral reasons in the actual world to pursue those strategies that minimise the possibilities of dystopian outcomes are very strong. This is in addition to epistemic reasons to expect large amounts of suffering as outlined above: humanity's track record and the tendency of evolutionary processes to produce large amounts of suffering, such as among wild animals.⁹

Finally, when making our assessments about the far future, we must take into account the possibility of being biased. The prospect of future dystopias is uncomfortable and unpleasant to think about. Most of us dread the possibility that our legacy in the universe could be a tragic one, and such a gloomy outlook does not resonate with favoured trends of techno-optimism or the heroic notion of saving humanity from extinction proposed by Bostrom (2003) and others. There are also other biases we should account for, such as normalcy bias, also known as the ostrich effect, which has mostly been discussed as a reason people fail to prepare for pandemics, extreme weather events, earthquakes, or other disasters because they expect their own life to continue as normal (Moon, Hwang, & Chung, 2019). Not all individuals may be biased in the same direction in these matters, but those who are may therefore irrationally discount the likelihood of astronomically bad future outcomes. An emphasis on moral circle expansion may be a good strategy to compensate for this bias.

5.3. Artificial sentience

Bostrom (2002) introduced the term 'existential risk', or *x-risk*, to refer to the possibility of 'an adverse outcome [that] would either annihilate Earth-originating intelligent life or permanently and drastically curtail its potential'. Work in this field has focused on the annihilation possibility, especially extinction due to artificial superintelligence (Bostrom, 2014). However, there is also the threat of *suffering risks* or *s-risks*, *x-risks* that entail not annihilation but 'suffering on an astronomical scale' (Daniel, 2017). These risks have only recently begun to receive substantial attention despite the vast moral stakes involved (Tomasik, 2015a). Given humanity's track record, we should at least consider the possibility of future atrocities and not rely on the inevitability of highly positive scenarios for the

⁸ A negative-leaning axiology need not be that simple. For instance, Benatar (2006) famously claims that the absence of pain is impersonally good, even though the absence of pleasure is not impersonally bad. Even if one believes that units of pleasure and pain have similar moral weight when present, the asymmetrical value of their absence gives us additional reasons to prevent the presence of suffering than to ensure the presence of pleasure.

⁹ Some may argue that a focus on avoiding dystopian outcomes neglects the possibility that positive experiences could be far better than what we can currently imagine. We should consider this possibility, but it appears at least as likely that we will transcend *negative* earthly experiences and endure much greater suffering than we can currently imagine.

continuation of human existence.

S-risks could develop in a similar way to the factory farming of animals. Suppose that humanity were to exploit some sort of artificial being on an astronomical scale for the production of a product or service valued by humans or human descendants, such as galactic-scale construction projects. Other examples include recreation, entertainment, or scientific experimentation motivated by curiosity or technological development. Perhaps the most hazardous s-risk scenarios come from the creation of *dolorium*, computer hardware designed to produce maximum suffering, which could be far worse than readily imaginable negative outcomes and might be produced for antagonistic reasons of revenge, threats, sadism, or blackmail.

While moral circle expansion is the s-risk reduction approach favoured by the nonprofit think tank Sentience Institute (Reese, 2018), it is not the only way to prevent s-risks. Some argue that artificial intelligence, in particular an artificial general intelligence, is one of the most likely ways that s-risks could occur (Tomasik, 2015a). Therefore, one could work to prevent s-risks by technical research on artificial intelligence itself, such as building mathematical and computer science tools that could be incorporated into an advanced AI system to reduce the chances of astronomical harm; this is a promising approach favoured by the Center on Long-Term Risk (Oesterheld, 2016).

6. Conclusion

We believe that those concerned with doing the most good, or even those simply trying to avoid serious wrongs, should prioritise the strategy of expanding our circle of moral concern to include all sentient beings, both in ethical theory and practice.

What does this entail? We suggest that initial practical efforts focus on increasing concern for those individuals presently at the frontiers of the moral circle such as the global poor, nonhuman animals, and future generations affected by climate change. In particular, increased concern for farmed animals may be the gateway to consideration for future sentient beings who may exist in astronomical numbers and who may be very dissimilar to humans. Two other promising approaches are research—developing better answers to the questions outlined above as well as more particular questions of different expansion strategies (see, for example, Mohorčič & Reese, 2019; Reese, 2020)—and movement-building or field-building—increasing the number of people focused on moral circle expansion and helping them develop their interests and skills in this field, either in research or practice.

Institutions, both public and private, should also take steps to consider wider scopes of beings in their constitutions, charters, and policies. Moral change-makers and researchers should keep an eye out for leverage points of influence on the far future. For example, given that the takeoff of artificial general intelligence could heavily shape humanity's trajectory, it is particularly important that relevant decision-makers take into account the interests of all sentient beings.

Finally, the consideration of the far future suggests that moral change-makers and researchers should focus on actions that are most likely to effect long-term moral circle expansion, for instance, by campaigning for the legal personhood of animals, which may directly affect only a small number of beings in the short run but could facilitate the legal personhood of artificial sentience with a similar socio-political status, or emphasising the interconnectedness of various inclusionary issues rather than discussing each of them in isolation.

Funding

No specific financial support was received for this article.

Availability of data and material

Not applicable

Code availability

Not applicable

Declaration of Competing Interest

The authors declare that they have no competing interests.

Acknowledgements

Thanks to Kelly Anthis for contributions to the initial conceptualization of moral circle expansion, Daniela R. Waldhorn for feedback on an early draft of this paper, and Tyler John for later stage feedback.

References

- Ahvenharju, S., Minkinen, M., & Lalot, F. (2018). The five dimensions of futures consciousness. *Futures*, 104(December), 1–13. <https://doi.org/10.1016/j.futures.2018.06.010>.
- Anthis, J. R. (2018). *Animals and the far future* (EAGxAustralia). <https://www.youtube.com/watch?v=NTV81NZSuKw>.
- Anthis, J. R. (2019). *US factory farming estimates*. Sentience Institute. <http://www.sentienceinstitute.org/us-factory-farming-estimates>.

- Anthis, K., & Anthis, J. R. (2019). *Global farmed & factory farmed animals estimates*. Sentience Institute. <http://www.sentienceinstitute.org/global-animal-farming-estimates>.
- Beard, S., Rowe, T., & James, F. (2020). An analysis and evaluation of methods currently used to quantify the likelihood of existential hazards. *Futures*, 115(January), 102469. <https://doi.org/10.1016/j.futures.2019.102469>.
- Belmonte, K. C. (2007). *William Wilberforce: A Hero for humanity*. Zondervan: Grand Rapids, Mich.
- Benatar, D. (2006). *Better never to have been: The harm of coming into existence*. New York: Clarendon Press.
- Bentham, J. (1789). *An introduction to the principles of morals and legislation*.
- Bentham, J., & Crompton, L. (1978). Offences against one's self: Paederasty (Part 1). *Journal of Homosexuality*, 3(4), 389–406. https://doi.org/10.1300/J082v03n04_07.
- Bostrom, N. (2002). Existential risks: Analyzing human extinction scenarios and related hazards. *Journal of Evolution and Technology*, 9. <https://ora.ox.ac.uk/objects/uuid:827452c3-fc6a-41b8-86b0-407293e6617c>.
- Bostrom, N. (2003). Astronomical waste: The opportunity cost of delayed technological development. *Utilitas*, 15(3), 308–314. <https://doi.org/10.1017/S095382080004076>.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies* (First edition). Oxford: Oxford University Press.
- Buttazzo, G. (2001). Artificial consciousness: Utopia or real possibility? *Computer*, 34(7), 24–30. <https://doi.org/10.1109/2.933500>.
- Caviola, L., Everett, J. A. C., & Faber, N. S. (2018). The moral standing of animals: Towards a psychology of speciesism. *Journal of Personality and Social Psychology*, (March) <https://doi.org/10.1037/pspp0000182>.
- Christiano, P. (2013). *Against moral advocacy*. Rational Altruist (blog). <https://rationalaltruist.com/2013/06/13/against-moral-advocacy/>.
- Cochrane, A. (2018). *Sentientist politics: A theory of global inter-species justice* (First edition). Oxford, United Kingdom: Oxford University Press.
- Costello, K., & Hodson, G. (2010). Exploring the roots of dehumanization: The role of animal–Human similarity in promoting immigrant humanization. *Group Processes & Intergroup Relations*, 13(1), 3–22. <https://doi.org/10.1177/1368430209347725>.
- Costello, K., & Hodson, G. (2014). Explaining dehumanization among children: The interspecies model of prejudice. *The British Journal of Social Psychology*, 53(1), 175–197. <https://doi.org/10.1111/bjso.12016>.
- Crimston, D., Bain, P. G., Hornsey, M. J., & Bastian, B. (2016). Moral expansiveness: Examining variability in the extension of the moral world. *Journal of Personality and Social Psychology*, 111(4), 636–653. <https://doi.org/10.1037/pspp0000086>.
- Daniel, M. (2017). *S-risks: Why they are the worst existential risks, and how to prevent them* (EAG Boston 2017). Foundational Research Institute (blog). <https://foundational-research.org/s-risks-talk-eag-boston-2017/>.
- Pierce, D. (1995). *The hedonistic imperative*. <https://www.hedweb.com/>.
- Dhont, K., & Hodson, G. (2014). Why do right-wing adherents engage in more animal exploitation and meat consumption? *Personality and Individual Differences*, 64 (July), 12–17. <https://doi.org/10.1016/j.paid.2014.02.002>.
- Dhont, K., Hodson, G., Costello, K., & MacInnis, C. C. (2014). Social dominance orientation connects prejudicial human–Human and human–Animal relations. *Personality and Individual Differences*, 61–62(April), 105–108. <https://doi.org/10.1016/j.paid.2013.12.020>.
- Dhont, K., Hodson, G., & Leite, A. C. (2016). Common ideological roots of speciesism and generalized ethnic prejudice: The social dominance human–animal relations model (SD-HARM): The social dominance human–animal relations model. *European Journal of Personality*, 30(6), 507–522. <https://doi.org/10.1002/per.2069>.
- Donaldson, S., & Kymlicka, W. (2011). *Zoopolis: A political theory of animal rights*. Oxford: Oxford University Press.
- Dorsser, C., Walker, W. E., Taneja, P., & Marchau, V. A. W. J. (2018). Improving the link between the futures field and policymaking. *Futures*, 104(December), 75–84. <https://doi.org/10.1016/j.futures.2018.05.004>.
- Drescher, S. (1977). *Econocide: British slavery in the era of abolition*. Pittsburgh: University of Pittsburgh Press.
- Easterbrook, G. (2018). *It's better than it looks* (First edition). New York: PublicAffairs.
- Faria, C. (2014). Equality, Priority and Nonhuman Animals. *Dilemata*, 6(4), 225–236.
- Faria, C. (2016). Why we should not postpone awareness of wild animal suffering. *Animal Sentience*, 1(7), 14.
- Faria, C., & Paez, E. (2015). Animals in need: The problem of wild animal suffering and intervention in nature. *Relations: Beyond Anthropocentrism*, 3(7).
- Graça, J., Calheiros, M. M., Oliveira, A., & Milfont, T. L. (2018). Why are women less likely to support animal exploitation than men? The mediating roles of social dominance orientation and empathy. *Personality and Individual Differences*, 129(July), 66–69. <https://doi.org/10.1016/j.paid.2018.03.007>.
- Greaves, H., & MacAskill, W. (2019). *The case for strong longtermism*. Global Priorities Institute. No.7-2019.
- Griffin, D. R. (1981). *The question of animal awareness*. Los Altos: William Kaufmann.
- Hodson, G., MacInnis, C. C., & Costello, K. (2013). Overvaluing 'humanness' as an aggravator of intergroup prejudices and discrimination. In P. G. Bain (Ed.), *Humanness and dehumanization*. New York: Psychology Press.
- Horta, O. (2010). Debunking the idyllic view of natural processes: Population dynamics and suffering in the wild. *Telos: Revista Iberoamericana de Estudios Utilitaristas*, 17(1), 73–90.
- Horta, O. (2018). Moral considerability and the argument from relevance. *Journal of Agricultural & Environmental Ethics*, 31(3), 369–388. <https://doi.org/10.1007/s10806-018-9730-y>.
- HSA. (2018). *Humane slaughter of finfish farmed around the world*. Humane Slaughter Association. <https://www.hsa.org.uk/downloads/hsafishslaughterreportfeb2018.pdf>.
- HSUS. (2013). *An HSUS report: The welfare of animals in the chicken industry*. Humane Society of the United States. <https://www.humanesociety.org/sites/default/files/docs/hsus-report-welfare-chicken-industry.pdf>.
- Judge, M., & Wilson, M. S. (2015). Vegetarian utopias: Visions of dietary patterns in future societies and support for social change. *Futures*, 71(August), 57–69. <https://doi.org/10.1016/j.futures.2015.07.005>.
- Kagan, S. (2019). *How to count animals, more or less*. Oxford: Oxford University Press.
- Kareiva, P., & Carranza, V. (2018). Existential risk due to ecosystem collapse: Nature strikes back. *Futures*, 102(September), 39–50. <https://doi.org/10.1016/j.futures.2018.01.001>.
- Korsgaard, C. M. (2018). *Fellow creatures: Our obligations to the other animals* (First edition). Oxford: Oxford University Press. Uehiro Series in Practical Ethics.
- Kovic, M. (2021). Risks of space colonization. *Futures*, 126(February), Article 102638. <https://doi.org/10.1016/j.futures.2020.102638>.
- Lecky, W. E. H. (1869). *History of European morals from Augustus to Charlemagne, 2 vols*. New York: D. Appleton and Company.
- Leopold, A. (1949). *A sand county almanac and sketches Here and there*. New York: Oxford University Press.
- Lippert-Rasmussen, K. (2013). *Born free and equal? A philosophical inquiry into the nature of discrimination*. Oxford: Oxford University Press.
- Liu, H.-Y., Lauta, K. C., & Maas, M. M. (2018). Governing boring apocalypses: a new typology of existential vulnerabilities and exposures for existential risk research. *Futures*, 102(September), 6–19. <https://doi.org/10.1016/j.futures.2018.04.009>.
- Low, P. (2012). *Cambridge declaration on consciousness*. <http://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>.
- MacAskill, W., & Ord, T. (2018). Why maximize expected choice-worthiness? *Nóus*, (July) <https://doi.org/10.1111/nous.12264>.
- McMahan, J. (2015). The moral problem of predation. In A. Chignell, T. Cuneo, & M. C. Halteman (Eds.), *Philosophy comes to dinner*. <https://doi.org/10.4324/9780203154410-23>.
- Midgley, M. (1983). *Animals and why they matter*. Athens: University of Georgia Press.
- Mill, J. S. (1859). *On liberty*.
- Mohorich, J., & Anthis, J. R. (2018). *What can the adoption of GM foods teach us about the adoption of other food technologies?* Sentience Institute. <https://www.sentienceinstitute.org/gm-foods>.
- Mohorich, J., & Reese, J. (2019). Cell-cultured meat: Lessons from GMO adoption and resistance. *Appetite*, 143(December), 104408. <https://doi.org/10.1016/j.appet.2019.104408>.

- Moon, J.-W., Hwang, H., & Chung, J.-B. (2019). Factors affecting awareness of preparedness after moderate earthquakes: An analysis of the Pohang earthquake in Korea. *Disaster Prevention and Management: An International Journal*, 29(3), 405–420. <https://doi.org/10.1108/DPM-07-2019-0209>.
- Moynihan, T. (2020). Existential risk and human extinction: An intellectual history. *Futures*, 116(February), Article 102495. <https://doi.org/10.1016/j.futures.2019.102495>.
- MSPCA. (2015). *George Thorndike Angell*. MSPCA-Angell (blog). <https://www.mspca.org/who-we-are/history/george-thorndike-angell/>.
- NCC. (2019). *U.S. Broiler performance*. The National Chicken Council (blog). <https://www.nationalchickencouncil.org/about-the-industry/statistics/u-s-broiler-performance/>.
- Ng, Y.-K. (2020). Effective altruism despite the second-best challenge: Should indirect effects be taken into account for policies for a better future? *Futures*, 121 (August), 102568. <https://doi.org/10.1016/j.futures.2020.102568>.
- Nussbaum, M. C. (2006). The moral status of animals. *The Chronicle of Higher Education*, 52(22), B6–8.
- Oesterheld, C. (2016). *Backup utility functions: A fail-safe AI technique*. November 21, 2016. Foundational Research Institute (blog) <https://foundational-research.org/backup-utility-functions-fail-safe-ai-technique/>.
- Ord, T. (2020). *The precipice: Existential risk and the future of humanity*. New York: Hachette Books.
- Paez, E. (2019). On the importance of species for rule-consequentialism: A reply to Galvão. *Utilitas*, 31(2), 179–190. <https://doi.org/10.1017/S0953820818000183>.
- Paez, E. (2020). A Kantian ethics of paradise engineering. *Analysis*, 80(2), 283–293. <https://doi.org/10.1093/analys/anz077>.
- Paez, E. (2021). A republic for all sentients: Social freedom without free will. *Pacific Philosophy Quarterly*. <https://doi.org/10.1111/papq.12351>.
- Palmer, C. (2010). *Animal ethics in context*. New York: Columbia University Press.
- Parfit, D. (1992). *Reasons and persons*. Repr. Oxford: Clarendon.
- Pinker, S. (2012). *The better angels of our nature*. New York Toronto London: Penguin Books.
- Pinker, S. (2018). *Enlightenment now*. New York, New York: Viking, an imprint of Penguin Random House LLC.
- Ramelli, I. (2009). *Hierocles the stoic: Elements of ethics, fragments and excerpts*, v. 28. Atlanta: Society of Biblical Literature. Translated by David Konstan. Society of Biblical Literature Writings from the Greco-Roman World.
- Reese, J. (2017). *Survey of US attitudes towards animal farming and animal-free food*. Sentience Institute. <http://www.sentienceinstitute.org/animal-farming-attitudes-survey-2017>.
- Reese, J. (2018). *The end of animal farming: How scientists, entrepreneurs, and activists are building an animal-free food system*. Boston: Beacon Press.
- Reese, J. (2020). Institutional change and the limitations of consumer activism. *Palgrave Communications*, 6(1). <https://doi.org/10.1057/s41599-020-0405-8>.
- Reinhart, R. J. (2018). *Snapshot: Few Americans vegetarian or vegan*. Gallup. <https://news.gallup.com/poll/238328/snapshot-few-americans-vegetarian-vegan.aspx>.
- Rodman, J. (1979). Animal justice: The counter-revolution in natural right and law. *Inquiry*, 22(1–4), 3–22. <https://doi.org/10.1080/00201747908601864>.
- Sidanius, J., & Pratto, F. (1999). *Social dominance*. Cambridge: Cambridge Univ. Press.
- Singer, P. (1981). *The expanding circle: Ethics and socio-biology*. New York: Farrar Straus & Giroux.
- Singer, P. (2015). *The most good you can do: How effective altruism is changing ideas about living ethically. Castle lectures in ethics, politics, and economics*. New Haven ; London: Yale University Press.
- Smith, A., & Reese, J. (2016). An empirical perspective on animal advocacy. *Animal Sentience*, 1(7). <https://animalstudiesrepository.org/animsent/vol1/iss7/12>.
- Taylor, P. W. (1981). The ethics of respect for nature. *Environmental Ethics*, 3(3), 197–218. <https://doi.org/10.5840/enviroethics19813321>.
- Tomasik, B. (2009). *How many wild animals are there?* Reducing Suffering. <https://reducing-suffering.org/how-many-wild-animals-are-there/>.
- Tomasik, B. (2015a). *Risks of astronomical future suffering*. Foundational Research Institute (blog). <https://foundational-research.org/risks-of-astronomical-future-suffering/>.
- Tomasik, B. (2015b). The importance of wild-animal suffering. *Relations. Beyond Anthropocentrism*, 3(2), 133–152. <https://doi.org/10.7358/rela-2015-002-toma>.
- Williams, E. G. (2015). The possibility of an ongoing moral catastrophe. *Ethical Theory and Moral Practice*, 18(5), 971–982. <https://doi.org/10.1007/s10677-015-9567-7>.
- Wright, D. W. M. (2018). Cloning animals for tourism in the year 2070. *Futures*, 95(January), 58–75. <https://doi.org/10.1016/j.futures.2017.10.002>.